

QNX V6.5 ON ARM

© Copyright Dedicated Systems Experts NV. All rights reserved, no part of the contents of this document may be reproduced or transmitted in any form or by any means without the written permission of Dedicated Systems Experts NV, Diepenbeemd 5, B-1650 Beersel, Belgium.

Authors: Luc Perneel (1, 2), Hasan Fayyad-Kazan (2) and Martin Timmerman (1, 2, 3)

1: Dedicated Systems Experts, 2: VUB-Brussels, 3: RMA-Brussels

Disclaimer

Although all care has been taken to obtain correct information and accurate test results, Dedicated Systems Experts, VUB-Brussels, RMA-Brussels and the authors cannot be liable for any incidental or consequential damages (including damages for loss of business, profits or the like) arising out of the use of the information provided in this report, even if these organisations and authors have been advised of the possibility of such damages.

<http://www.dedicated-systems.com>

E-mail: info@dedicated-systems.com

EVALUATION REPORT LICENSE

This is a legal agreement between you (the downloader of this document) and/or your company and the company DEDICATED SYSTEMS EXPERTS NV, Diepenbeemd 5, B-1650 Beersel, Belgium.
It is not possible to download this document without registering and accepting this agreement on-line.

1. **GRANT.** Subject to the provisions contained herein, Dedicated Systems Experts hereby grants you a non-exclusive license to use its accompanying proprietary evaluation report for projects where you or your company are involved as major contractor or subcontractor. You are not entitled to support or telephone assistance in connection with this license.
2. **PRODUCT.** Dedicated Systems Experts shall furnish the evaluation report to you electronically via Internet. This license does not grant you any right to any enhancement or update to the document.
3. **TITLE.** Title, ownership rights, and intellectual property rights in and to the document shall remain in Dedicated Systems Experts and/or its suppliers or evaluated product manufacturers. The copyright laws of Belgium and all international copyright treaties protect the documents.
4. **CONTENT.** Title, ownership rights, and an intellectual property right in and to the content accessed through the document is the property of the applicable content owner and may be protected by applicable copyright or other law. This License gives you no rights to such content.
5. **YOU CANNOT:**
 - You cannot, make (or allow anyone else make) copies, whether digital, printed, photographic or others, except for backup reasons. The number of copies should be limited to 2. The copies should be exact replicates of the original (in paper or electronic format) with all copyright notices and logos.
 - You cannot, place (or allow anyone else place) the evaluation report on an electronic board or other form of on line service without authorisation.
6. **INDEMNIFICATION.** You agree to indemnify and hold harmless Dedicated Systems Experts against any damages or liability of any kind arising from any use of this product other than the permitted uses specified in this agreement.
7. **DISCLAIMER OF WARRANTY.** All documents published by Dedicated Systems Experts on the World Wide Web Server or by any other means are provided "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. This disclaimer of warranty constitutes an essential part of the agreement.
8. **LIMITATION OF LIABILITY.** Neither Dedicated Systems Experts nor any of its directors, employees, partners or agents shall, under any circumstances, be liable to any person for any special, incidental, indirect or consequential damages, including, without limitation, damages resulting from use of OR RELIANCE ON the INFORMATION presented, loss of profits or revenues or costs of replacement goods, even if informed in advance of the possibility of such damages.
9. **ACCURACY OF INFORMATION.** Every effort has been made to ensure the accuracy of the information presented herein. However Dedicated Systems Experts assumes no responsibility for the accuracy of the information. Product information is subject to change without notice. Changes, if any, will be incorporated in new editions of these publications. Dedicated Systems Experts may make improvements and/or changes in the products and/or the programs described in these publications at any time without notice. Mention of non-Dedicated Systems Experts products or services is for information purposes only and constitutes neither an endorsement nor a recommendation.
10. **JURISDICTION.** In case of any problems, the court of BRUSSELS-BELGIUM will have exclusive jurisdiction.

Agreed by downloading the document via the internet.

1	Document Intention.....	6
1.1	Purpose and scope	6
1.2	Document issue: the 2.9 framework	6
1.3	Related documents	7
2	Introduction	8
2.1	Overview	8
2.2	Evaluated (RTOS) product.....	8
2.2.1	Software	8
2.2.2	Hardware	8
3	Evaluation results summary.....	9
3.1	Positive points	9
3.2	Negative points	9
3.3	Ratings	9
4	Test Results.....	10
4.1	Calibration system test (CAL)	10
4.1.1	Tracing overhead (CAL-P-TRC).....	10
4.1.2	CPU power (CAL-P-CPU)	11
4.2	Clock tests (CLK)	12
4.2.1	Operating system clock setting (CLK-B-CFG)	12
4.2.2	Clock tick processing duration (CLK-P-DUR)	13
4.3	Thread tests (THR)	14
4.3.1	Thread creation behaviour (THR-B-NEW)	14
4.3.2	Round robin behaviour (THR-B-RR)	15
4.3.3	Thread switch latency between same priority threads (THR-P-SLS).....	15
4.3.4	Thread creation and deletion time (THR-P-NEW).....	18
4.4	Semaphore tests (SEM).....	23
4.4.1	Semaphore locking test mechanism (SEM-B-LCK)	23
4.4.2	Semaphore releasing mechanism (SEM-B-REL).....	23
4.4.3	Time needed to create and delete a semaphore (SEM-P-NEW).....	23
4.4.4	Test acquire-release timings: contention case (SEM-P-ARN)	26
4.4.5	Test acquire-release timings: contention case (SEM-P-ARC)	28
4.5	Mutex tests (MUT).....	31
4.5.1	Priority inversion avoidance mechanism (MUT-B-ARC)	31
4.5.2	Mutex acquire-release timings: contention case (MUT-P-ARC)	31
4.5.3	Mutex acquire-release timings: no-contention case (MUT-P-ARN)	33
4.6	Interrupt tests (IRQ)	35
4.6.1	Interrupt latency (IRQ_P_LAT).....	35
4.6.2	Interrupt dispatch latency (IRQ_P_DLT)	36
4.6.3	Interrupt to thread latency (IRQ_P_TLT).....	36
4.6.4	Maximum sustained interrupt frequency (IRQ_S_SUS).....	37
4.7	Memory tests.....	38

Doc: **EVA-2.9-TST-QNX-ARM-65**

Issue: **draft 3.10**

Date: **Sept 7, 2011**

4.7.1	Memory leak test (MEM_B_LEK)	38
5	Appendix A: Vendor comments	39
6	Appendix B: Acronyms	40

Doc: **EVA-2.9-TST-QNX-ARM-65**

 Issue: **draft 3.10**

 Date: **Sept 7, 2011**

DOCUMENT CHANGE LOG

Issue No.	Revised Issue Date	Para's / Pages Affected	Reason for Change
2.01	July 25, 2011	All	Initial draft – patch included
2.02	July 27, 2011	All	comments
3.00	July 28, 2011	All	vendor revision draft
3.1	September 7, 2011	All	Final version

1 Document Intention

1.1 Purpose and scope

This document presents the quantitative evaluation results of the QNX Neutrino operating system V6.5 employed on an ARM based platform, more specifically on the Beagle XM board.

The layout of this report follows the one depicted in “The OS evaluation template” [Doc. 4]. The test specifications can be found in “The evaluation test report definition.” [Doc. 3]. See section 1.3 of this document for more detailed references. These documents have to be seen as an integral part of this report!

Due to the tightly coupling between these documents, the framework version of “The evaluation test report definition.” has to match the framework version of this evaluation report (which is 2.9). More information about the documents and tests versions together with their corresponding relation between both can be found in “The evaluation framework”, see [Doc. 1] in section 1.3 of this document.

The generic test code used to perform these tests can be downloaded from our website by using the link in the related documents section.

1.2 Document issue: the 2.9 framework

This document shows the test results in the scope of the evaluation framework 2.9.

2 Introduction

This chapter talks about the OS that we are going to test and evaluate, and the hardware on which the OS under testing will be employed.

2.1 Overview

QNX Software Systems Ltd was founded in 1980 and has been always focused on delivering solutions for the embedded systems market.

One of the main differences between QNX and other RTOS is the fact that QNX is built around the POSIX API standard. This has its advantages as a lot of code for Linux based platforms can be compiled and run on QNX Neutrino. However, bear in mind that we are discussing a real-time operating system here.

QNX Neutrino is based on true microkernel architecture with message-based inter-process communication. For instance, drivers are just applications with special privileges, and as such they cannot crash the kernel. The concept of kernel modules which is the case in Linux is not needed here, which makes QNX Neutrino a very stable product.

Furthermore, QNX Neutrino was initially built-up as a multi-processor capable operating system (both SMP and AMP). Nowadays, this is a very important asset in today's multi- and many-core business.

2.2 Evaluated (RTOS) product

2.2.1 Software

The operating system that we are going to evaluate is the QNX NEUTRINO RTOS v6.5.0 including patch 2530, from QNX Software Systems Ltd.

2.2.2 Hardware

The hardware that was used for executing our tests for the QNX Neutrino RTOS was a Beagle-XM Board Rev C with following characteristics:

- based on the Texas Instruments DM3730 Digital Media Processor
- ARM Cortex A8 running at 1GHz
- L1 Cache: 32KB instruction and 32KB data cache
- L2 Cache: 64KB
- 512MB RAM at 166MHz

3 Evaluation results summary

Following is a summary of the results of evaluating the QNX NEUTRINO RTOS v6.5.0, from QNX Software Systems Ltd.

3.1 Positive points








- Excellent architecture for a robust and distributed system.
- Very fast and predictable performance.
- Large number of board support packages (BSP) and drivers (the source for most of them is available for public) which can be easily downloaded.
- The availability of documentation which can be considered more than the average.
- Efficient and user friendly Integrated Development Environment (IDE)

3.2 Negative points

- Not all code is available in source code. Customers can apply for source access.

3.3 Ratings

For a description of the ratings, see [Doc. 3].

RTOS Architecture	0		10
OS Documentation	0		10
OS Configuration	0		10
Internet Components	0		10
Development Tools	0		10
BSPs	0		10
Support	0		10

4 Test Results

Test Results 0  10

QNX 6.5 has very good performance characteristics which guarantee the real-time behaviour.

4.1 Calibration system test (CAL)

These tests are used to calibrate the tracing overhead compared with the processing power of the platform. This is important to understand the accuracy of the measurements done in scope of this report.

They are used also for measuring the processing power of the platform. This calibration permits comparison with the results on other platforms.

4.1.1 Tracing overhead (CAL-P-TRC)

As the Beagle board does not have any PCI support, we used the on-chip hardware timers for our measurements.

This test calibrates the tracing system overhead. It is related to the hardware more than to the OS, but it is needed to correct the measured times.

In the rest of the document, the tracing overhead is subtracted from the obtained results.

For tracing, an internal General Purpose (GP) timer running at 13MHz was used. Reading out these timers takes some overhead of course; however, there is not any jitter at all in the overhead of the trace which in turn does not generate much extra inaccuracies.

Although it is possible to let the general purpose timer run at a higher frequency, the clock that was attached to this GP timer is also distributed to other components on the chip as well. Therefore, we had to stick with the configuration that was used by the QNX BSP on this board.

In general, the results in this report are correct to +/- 0.2 μ seconds. Therefore the results shown in the tables are rounded to 0.1 microseconds.

4.1.1.1 Test results

Test	result
Number of samples	30000
Average tracing overhead	384 nsec
minimum tracing overhead	384 nsec
maximum tracing overhead	384 nsec

4.1.2 CPU power (CAL-P-CPU)

This test will calibrate the CPU performance and the memory bandwidth of the used platform. This test is measured in different situations: from the situation where code and data are cached, until the situation where neither code nor data are cached. With such different situation tests, the effects of the cache can be calculated.

We have been seriously reworking this test lately. The CPU test uses only one data address. The non-cached version is about 172KB in size (instructions), while the cached version uses a loop (a bit unrolled to have a small loop overhead but so it fits in the L1 I-cache and it uses only two data words). The instruction cache test is done twice.

4.1.2.1 Test results

The results for our standard platform (Pentium MMX 200 MHz) are shown below:

Test	no cache	cached	cache effect
CPU test: first load.	882 us		
CPU test: ICache effect	869 us	136.4 us	6.4
MEM write test	400 us	392.3 us	1.02
MEM read test	658 us	399.0 us	1.7
Average caching effect (CPU and MEM)			3.0

On the beagle platform, the same test gives following results:

Test	no cache	cached	cache effect
CPU test: first load.	399.5 us		
CPU test: ICache effect	230.1 us	31.5 us	7.3
MEM write test	29.1 us	21.1 us	1.4
MEM read test	44.2 us	21.2 us	2.1
Average caching effect (CPU and MEM)			3.6

The results show similar behaviour as on the Pentium MMX 200MHz:

- Instruction cache has most impact on performance (you need to get each instruction from RAM, while there are always less than one memory accesses for each instruction). Remark that this test is built on purpose and thus an extreme worst case that in practice will never occur (you will never have >100KB instructions without loops in real environments).
- Write can be postponed, so it has less impact
- As read is blocking, it has more impact than the writes.

If we compare the Beagle platform with our standard evaluation platform (Pentium MMX 200MHz) then we can conclude:

- The memory accesses on the Beagle platform are much faster than on the Pentium reference platform by a factor of 10 to 15. When cached, the difference is even more impressive
- Instruction loading outside the cache is not that fast compared with the Pentium. This will of course depend on the type of code. As the x86 has variable instruction lengths, most of the time the generated assembler code on the Pentium will be compacter and thus relatively faster to load.
- The performance differences between both boards depends largely on the type of the test and if cached or not. In the worst case, the 1GHz ARM is only marginally faster than the Pentium MMX 200MHz (un-cached instruction loading), but in the best case the difference is about a factor of 8 (cached writes). Thus, the relative performance between both platforms will enormously depend on the test scenario!

Because of the serious impact caused by not having your instructions in the cache, you should take extra safety margins in real-time behaviour (worst case is less predictable) on this platform.

Remark as well that the first load of code introduces some extra delay.

4.2 Clock tests (CLK)

The clock test measures the time that an operating system needs to handle its clock interrupt. On the tested platform, the clock tick interrupt is set on the highest hardware interrupt level, interrupting any other thread or interrupt handler.

4.2.1 Operating system clock setting (CLK-B-CFG)

This test is done in order to examine the setting of the clock tick period in the operating system. This test shows the default clock timing as they are set by the OS.

For this test, the `nanosleep()` POSIX function call is used. Following POSIX, the delay should be based on the clock tick. The “nanosleep” function always pauses for at least its specified time, but however it can take up to one clock tick more than its specified time until the process becomes run-able again.

As the OS cannot know when the next tick will occur, it will add the time of one tick and round it up to the higher tick. Using `nanosleep(0)` thus will wake up on next tick.

4.2.1.1 Test results

Test	result
Test succeeded	Yes
Tested clock period	1ms
Clock period adaptable	YES, using ClockPeriod().

4.2.2 Clock tick processing duration (CLK-P-DUR)

This test is done for examining the clock tick processing duration in the kernel. The test results are extremely important as the clock interrupt will disturb all the other performed measurements.

The bottom line of the figures in section 4.2.2.2 represents the normal loop time of the test if no clock interrupt occurs during the test loop. The upper line is generated by the samples when a clock interrupt occurred during the loop. The difference between the two lines is the clock tick processing duration.

In absolute values, the 2µsec clock tick impact is of course not that large on this fast platform.

Knowing that this platform has less processing power compared with the Atom Z540, QNX still manages to keep its clock tick interrupt as small as with the Atom!

4.2.2.1 Test results

Test	result
CLOCK_LOOP_COUNTER	5000
Normal busy loop time	37.5 µs
Busy loop time with clock interrupt	39.5 µs, worst case 46µs
Clock interrupt duration	2 µs to 6.5 µs

4.2.2.2 Diagrams

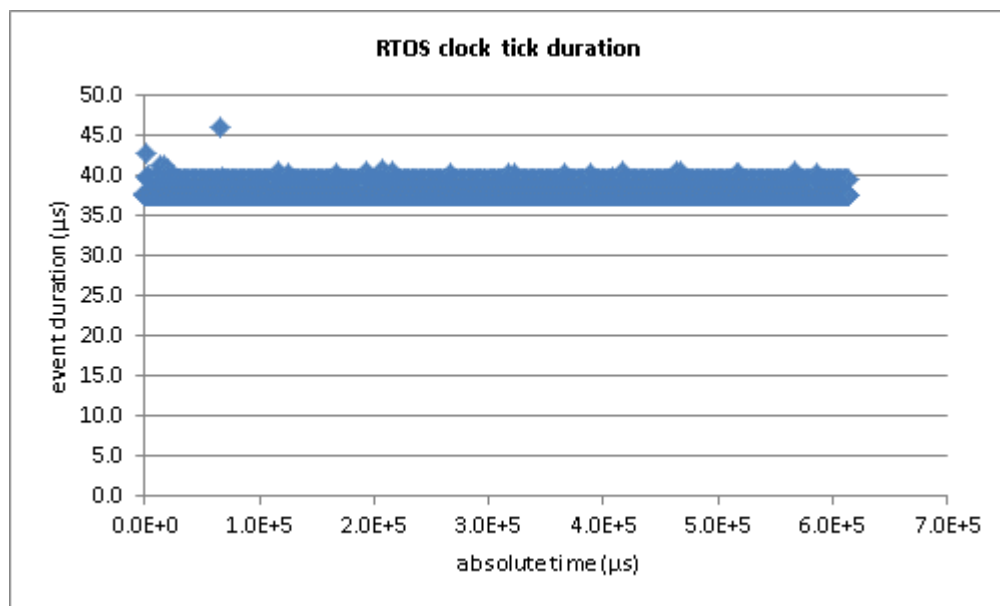


Figure 1: RTOS clock tick duration.

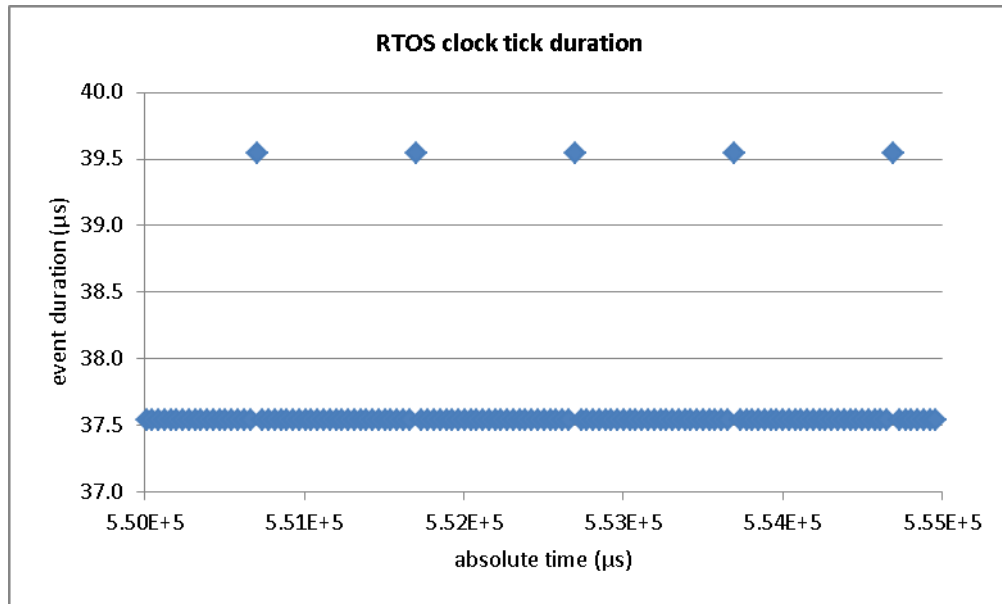


Figure 2: RTOS clock tick duration, detail.

4.3 Thread tests (THR)

These tests are used to measure the performance of the scheduler.

4.3.1 Thread creation behaviour (THR-B-NEW)

This test will examine the behaviour of creating threads. Does the operating system behave as it should be as long as it is considered being a real-time operating system? Following scenarios are tested:

- If a thread is created with a lower priority than the creating thread, then are we sure that it is not activated until the creating thread is finished?
- If a thread is created with the same priority as the creating thread, will it be put at the ready tail?
- When yielding after the creation in the above test, does the newly created thread becomes active?
- If a thread is created with a higher priority than the creating thread, is it then immediately activated?

This test succeeded without any problems.

4.3.1.1 Test results

Test	result
Test succeeded	YES
Lower priority not activated?	YES
Same priority at tail?	YES
Yielding works?	YES
Higher priority activated?	YES

4.3.2 Round robin behaviour (THR-B-RR)

This test checks if the scheduler uses a fair round robin mechanism when threads are having the same priority and all of them are in the ready-to-run state (and using the SCHED_RR scheduling policy)!

No problems were detected here. The round robin behaviour reschedules a thread each 4 clock ticks.

4.3.2.1 Test results

Test	result
Test succeeded	Yes
RR Time slice following this test	4 ms

4.3.3 Thread switch latency between same priority threads (THR-P-SLS)

This test measures the time to switch between threads of the same priority. Therefore, threads have to yield the processor voluntary for the other threads to use it.

In this test, we use the SCHED_FIFO policy; otherwise it would be possible that a round-robin clock event occurs between the yield and the trace, so that the thread activation is not seen in the trace.

This test was performed several times, and each time using a higher number of threads in order to generate the worst case behaviour. If more threads are active, then the caching effect will be obvious in a way that the thread context will not reside anymore in the L1 cache once we have enough threads and even not in the L2 cache later on. You can see this on the diagrams, as there is much more variation and the minimal value is about a factor three larger when running 1000 threads.

Further, you will clearly see the influence of clock interrupts (causing the higher values in the graphics).


Once there are enough running threads, the clock interrupt will always be un-cached and thus for the 1000 thread tests, the clock interrupts always generate a delay of approximately 4µs.

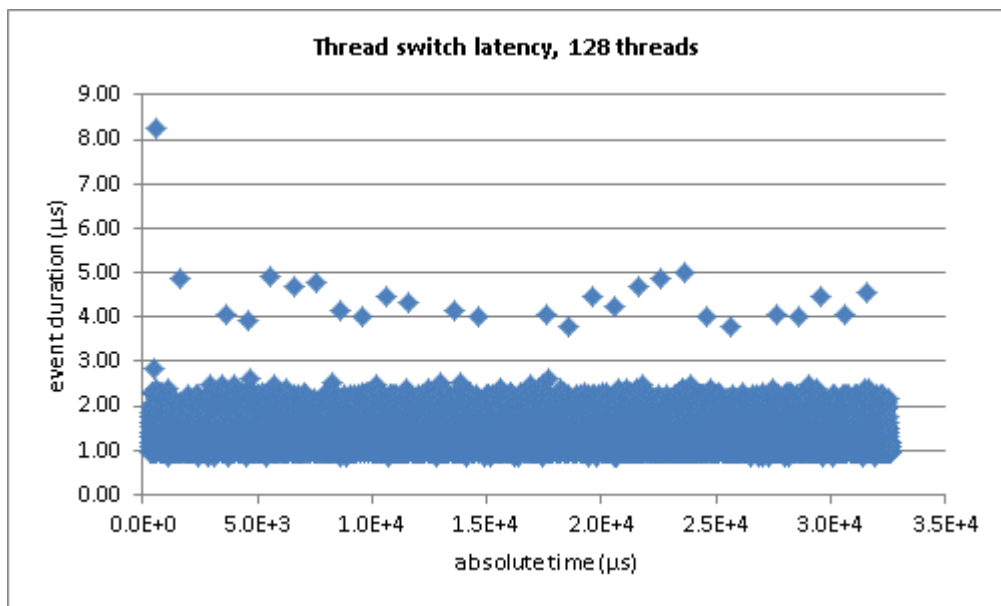
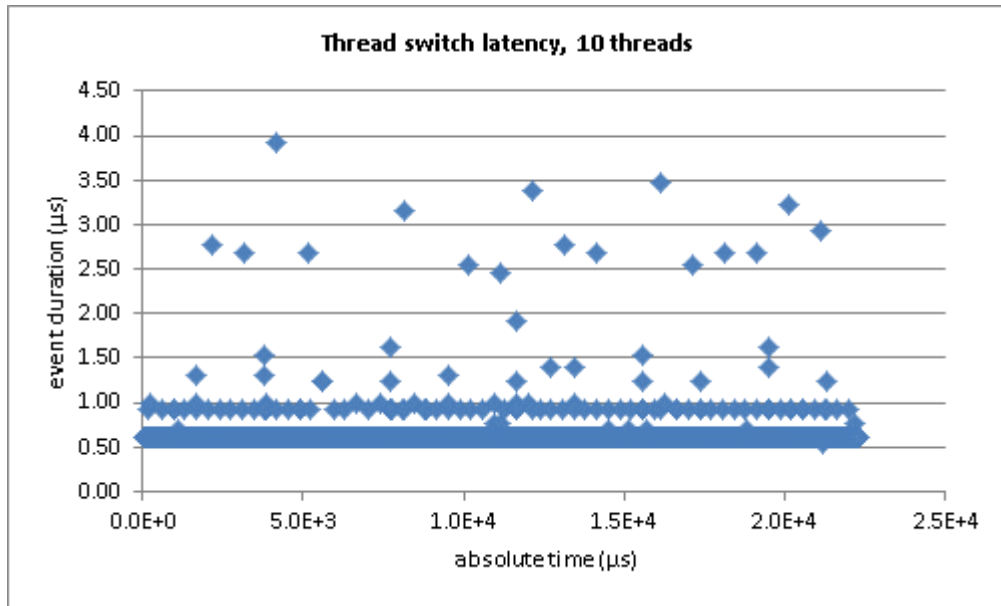
4.3.3.1 Test results

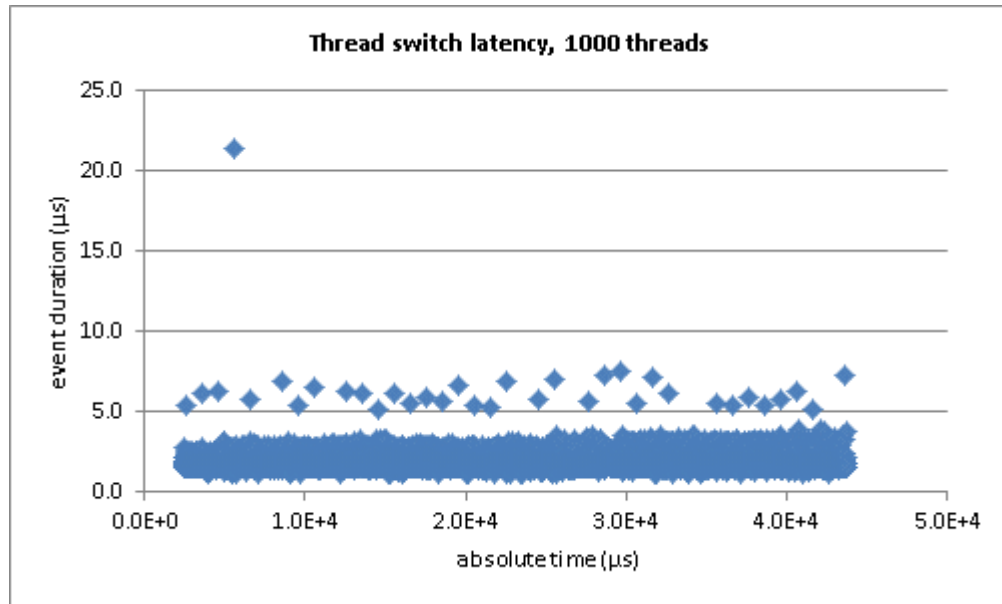
Test	result
Test succeeded	YES

Test	Sample qty	Avg	Max	Min
Thread switch latency, 2 threads	15999	0.6 μ s	6.3 μ s	0.6 μ s
Thread switch latency, 10 threads	15995	0.6 μ s	3.9 μ s	0.6 μ s
Thread switch latency, 128 threads	15936	1.2 μ s	8.2 μ s	0.8 μ s
Thread switch latency, 1000 threads	15500	1.9 μ s	21.5 μ s	1.1 μ s

4.3.3.2 Diagrams







4.3.4 Thread creation and deletion time (THR-P-NEW)

This test examines the time for creating a thread, and the time for deleting a thread in different scenarios:

- Scenario 1 “never run”: The created thread has a lower priority than the creating thread and is deleted before it has any chance to run. No thread switch occurs in this test.
- Scenario 2 “run and terminate”: The created thread has a higher priority than the creating thread and will be activated. The created thread immediately terminates itself (thread does nothing).
- Scenario 3 “run and block”: The same as the previous scenario (scenario 2: run and terminate), but the created thread does not terminate (it lowers its priority when it is activated).

In the scenarios where the thread actually runs (2, 3), the creation time is the duration from the system call creating the thread until the time when the created thread is activated. For the “never run” scenario, the creation time is the duration of the system call.

Clearly, the performance is much better than on our reference platform (Pentium 200MHz MMX) by a factor 4 or even more. Compared with the Z540 Atom platform, the performance is about half of the Atom.

At start-up, you can see sometimes a large spike due to caching.

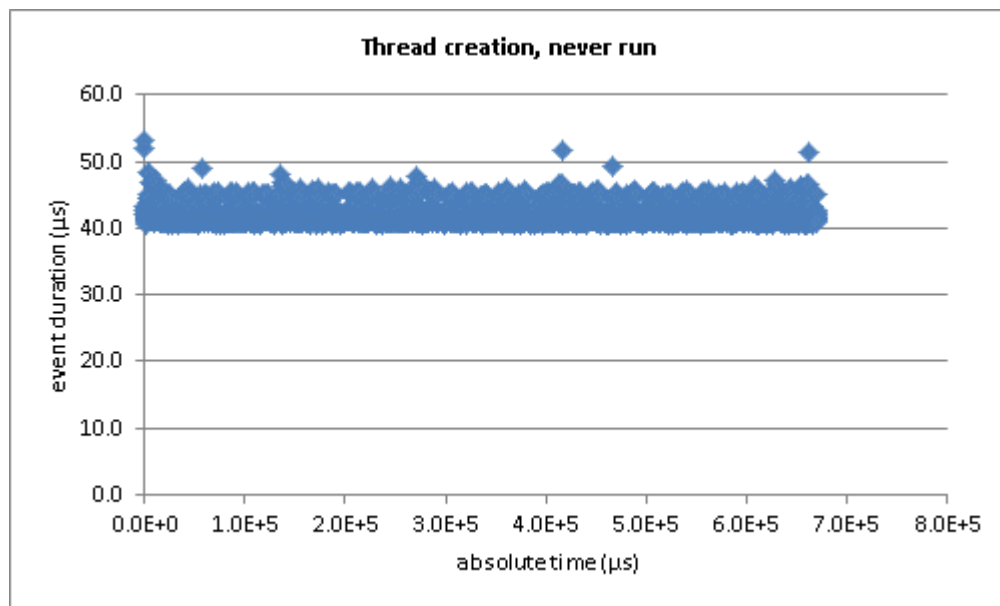
Comparing the Atom platform (which has 512KB L2 cache) with the Beagle XM platform (having only 64KB L2 cache), you can clearly see that the large cache on the Atom stabilizes the results, thus obtaining flat lines which is not the case on this ARM platform, which has some more variation caused by the cache misses.

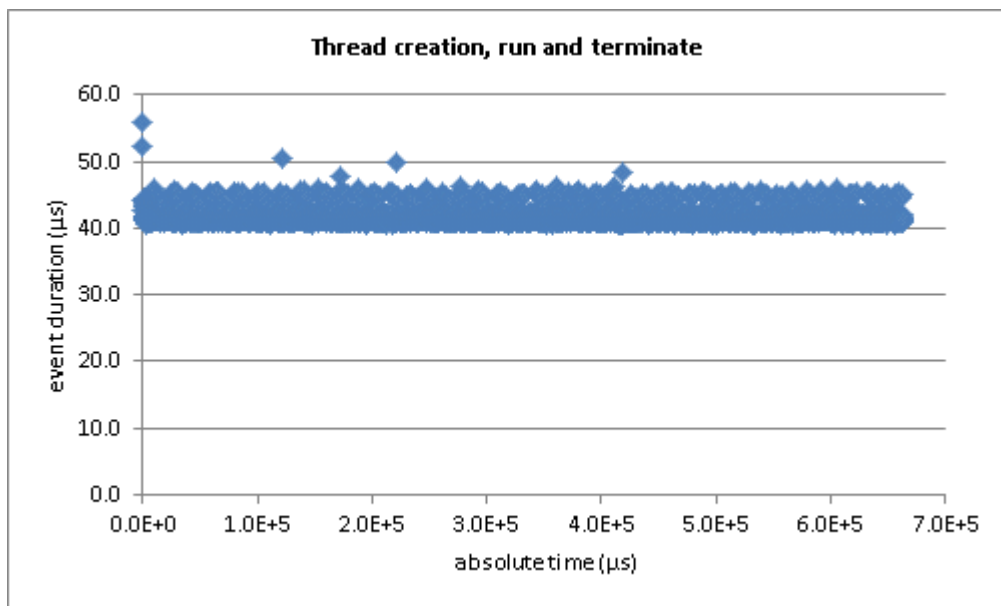
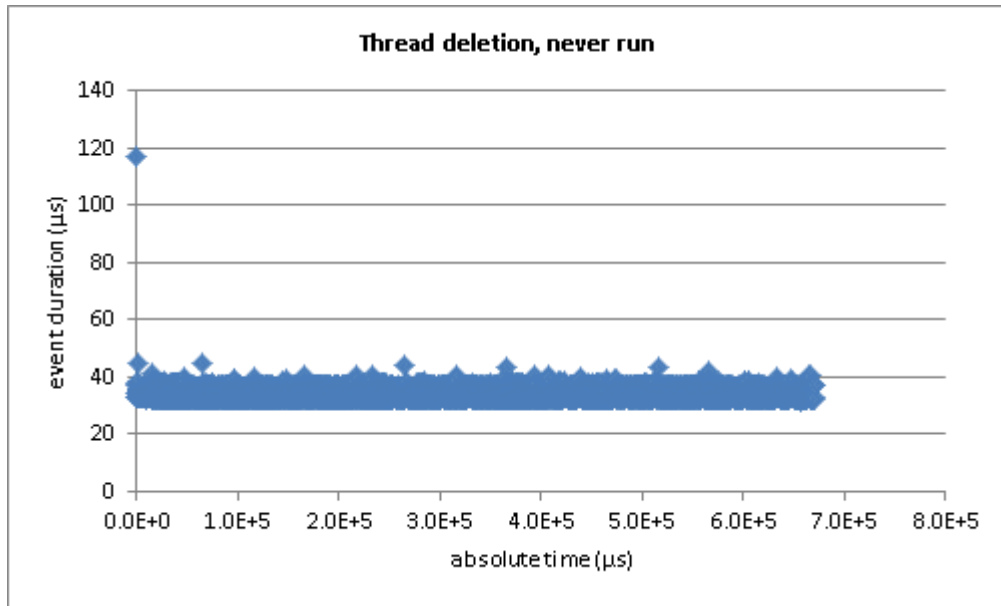
4.3.4.1 Test results

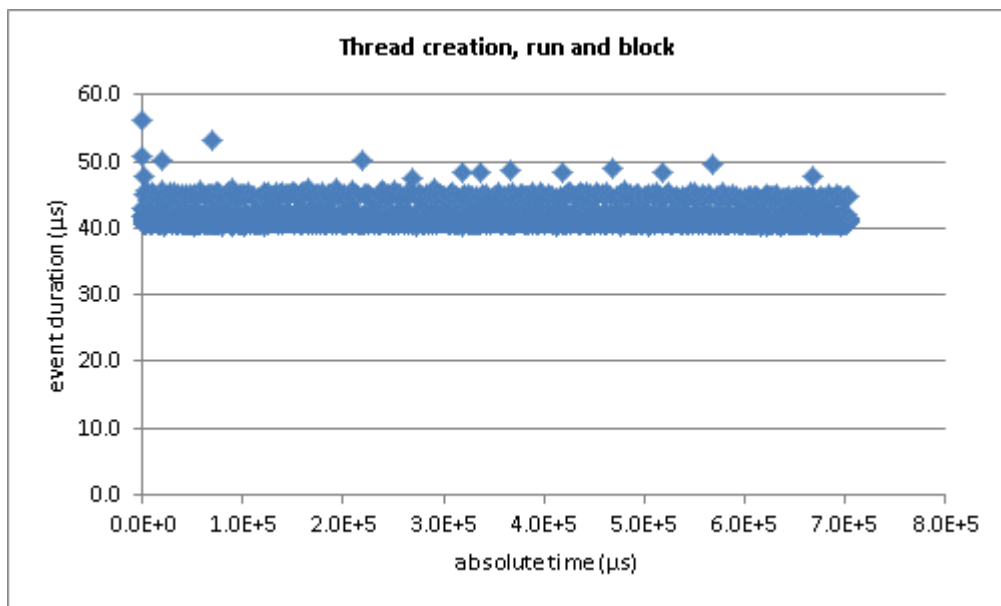
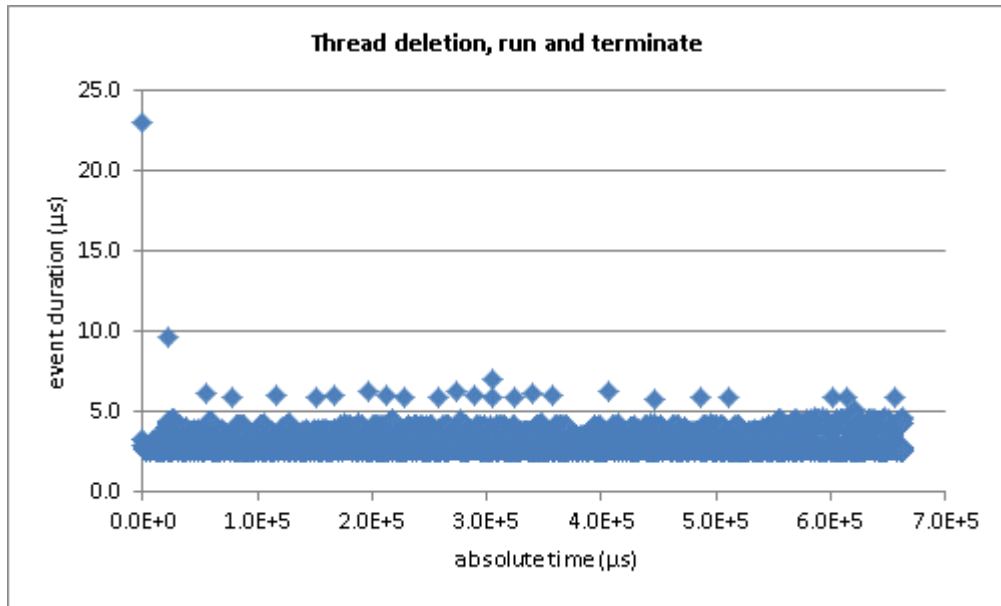
Test	result
Test succeeded	YES

Test	Sample qty	Avg	Max	Min
Thread creation, never run	8000	41.8 μ s	53.3 μ s	40.5 μ s
Thread deletion, never run	8000	34.4 μ s	117 μ s	31.5 μ s
Thread creation, run and terminate	8000	41.7 μ s	55.9 μ s	40.5 μ s
Thread deletion, run and terminate	8000	3.1 μ s	23.0 μ s	2.5 μ s
Thread creation, run and block	8000	41.5 μ s	56.3 μ s	40.2 μ s
Thread deletion, run and block	8000	35.4 μ s	116.5 μ s	32.5 μ s

4.3.4.2 Diagrams



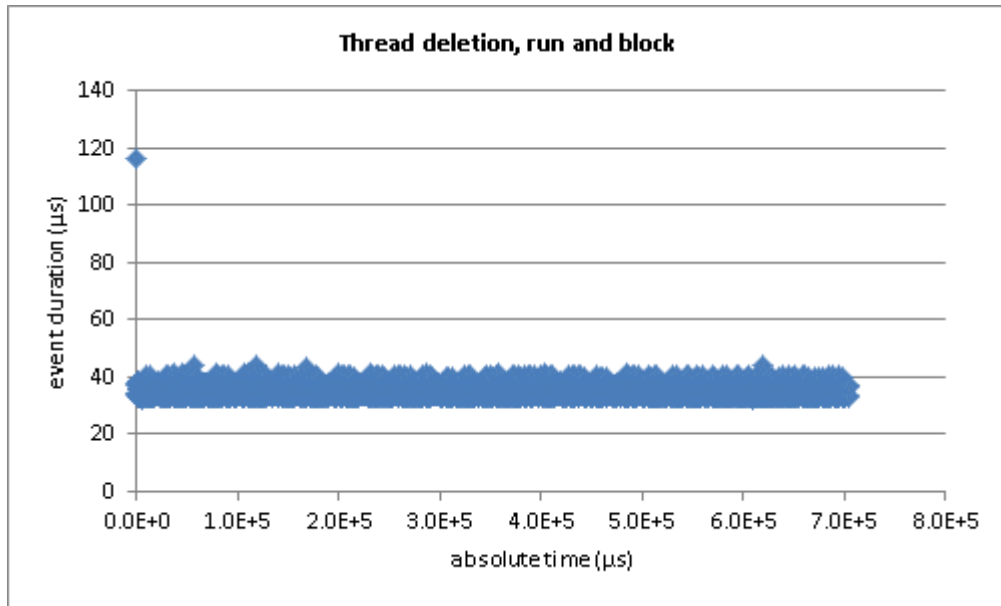




Doc: **EVA-2.9-TST-QNX-ARM-65**

Issue: **draft 3.10**

Date: **Sept 7, 2011**



4.4 Semaphore tests (SEM)

This test examines the performance and the behaviour of the counting semaphore. The counting semaphore is a system object that can be used to synchronize threads.

4.4.1 Semaphore locking test mechanism (SEM-B-LCK)

In this test, we will experiment if the counting semaphore locking mechanism works as it is expected to do. The $P()$ call should block only when the count is zero. The $V()$ call should increment the semaphore counter. In the case where the semaphore counter is zero, the $V()$ call should cause a rescheduling by the OS: indeed blocked threads may become active.

The semaphore behaves correctly as a protection mechanism.

4.4.1.1 Test results

Test	result
Test succeeded	YES
Maximum semaphore value?	Limited by the "int" type
Rescheduling on free?	OK

4.4.2 Semaphore releasing mechanism (SEM-B-REL)

This test verifies that the highest priority thread being blocked on a semaphore will be released by the release operation. This should be independent of the order of the acquisitions taking place.

QNX passed this test.

4.4.2.1 Test results

Test	result
Test succeeded	YES

4.4.3 Time needed to create and delete a semaphore (SEM-P-NEW)

This test is done to get an insight about the time needed to create a semaphore and the time to delete it. The deletion time is checked in two cases:

- The semaphore is used between the creation and deletion.
- The semaphore is NOT used between the creation and deletion.

Remark that although we do not use “named” semaphores, there seems to be a system call required to create/delete a semaphore.

At start up, there is a peak probably caused by caching.

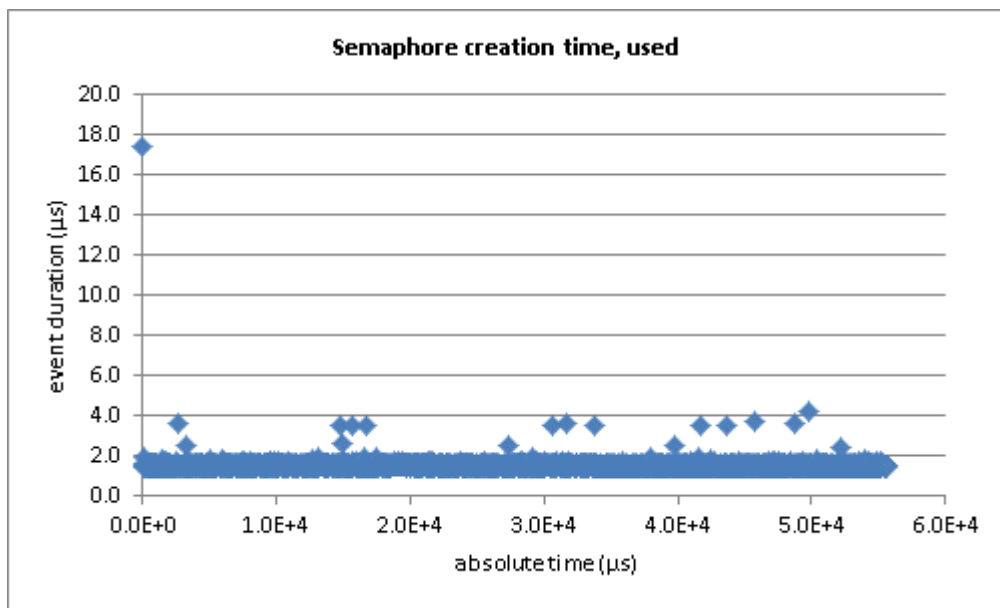
As usual, the clock tick interrupt is seen.


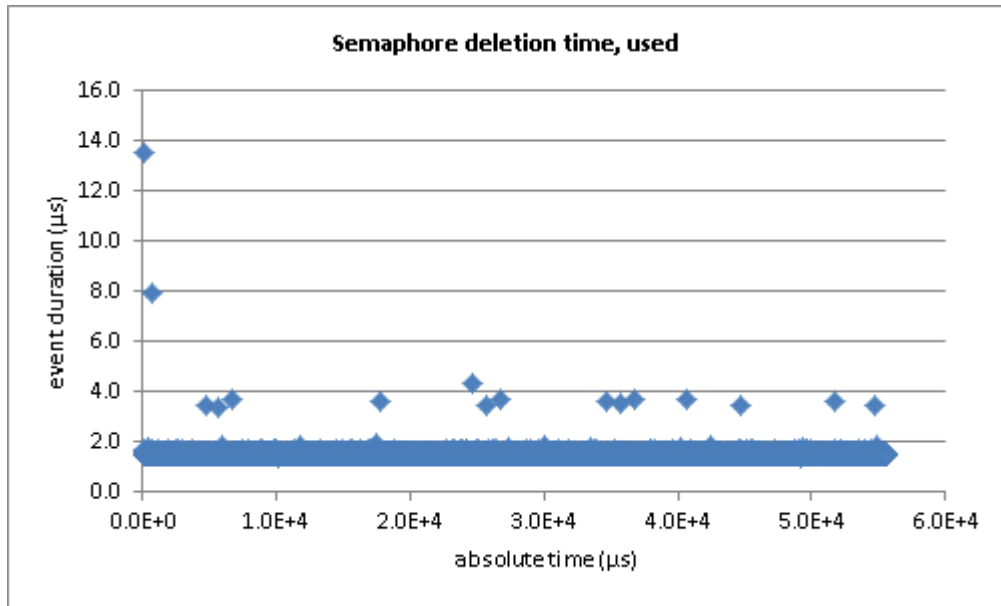
4.4.3.1 Test results

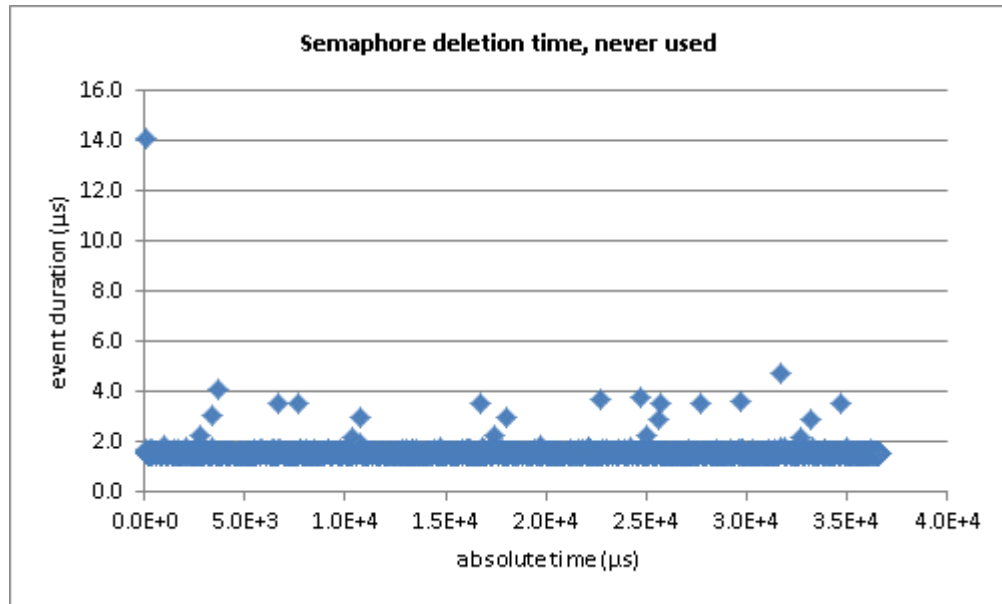
Test	result
Test succeeded	YES

Test	Sample qty	Avg	Max	Min
Semaphore creation time, used	8000	1.5 μ s	17.5 μ s	1.4 μ s
Semaphore deletion time, used	8000	1.5 μ s	13.5 μ s	1.4 μ s
Semaphore creation time, never used	8000	1.5 μ s	26.3 μ s	1.5 μ s
Semaphore deletion time, never used	8000	1.6 μ s	14.2 μ s	1.5 μ s

4.4.3.2 Diagrams







4.4.4 Test acquire-release timings: contention case (SEM-P-ARN)

Here we test the acquisition and release time in the non-contention case. As in this test case the semaphore does not neither block nor cause any rescheduling (thread switch), the duration of the call should be short.

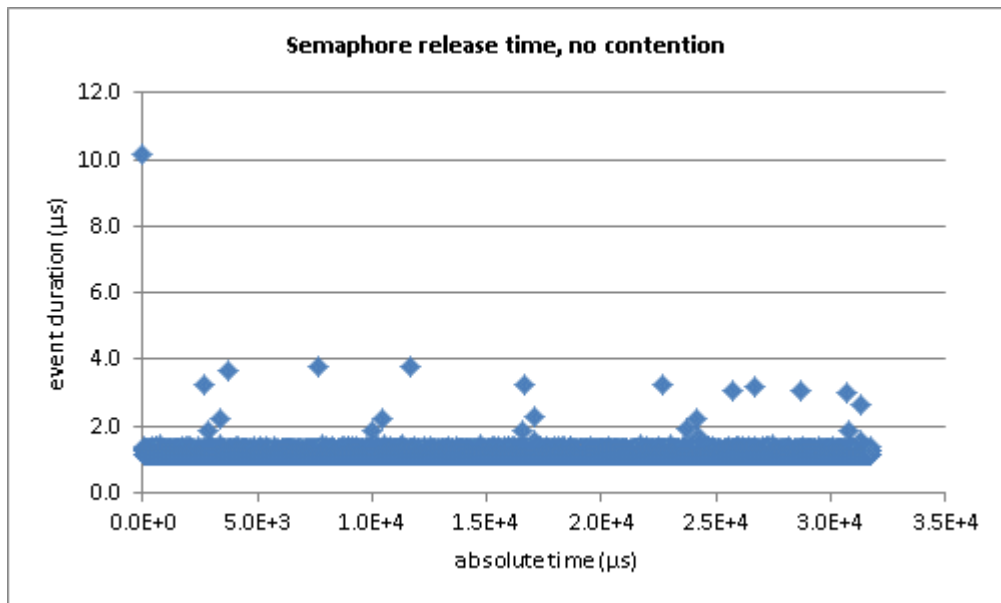
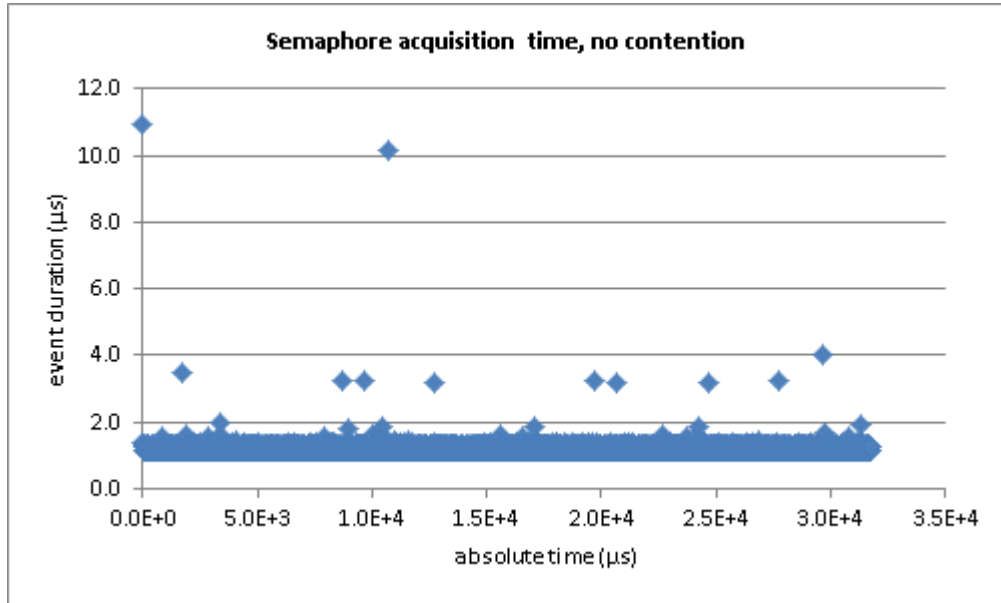
The clock tick is always present.

4.4.4.1 Test results

Test	result
Test succeeded	YES

Test	Sample qty	Avg	Max	Min
Semaphore acquisition time, no contention	8000	1.2 µs	10.1 µs	1.2 µs
Semaphore release time, no contention	8000	1.2 µs	10.2 µs	1.2 µs

4.4.4.2 Diagrams



4.4.5 Test acquire-release timings: contention case (SEM-P-ARC)

This is performed to test the time needed to acquire and release a semaphore, depending on the number of threads blocked on the semaphore. It measures the time in the contention case when the acquisition and release system call causes a rescheduling to occur.

The aim of this test is to verify whether the number of blocked threads has an impact on these timings or not. So this will answer the question: "how much time the operating system needs to find out the next thread to schedule".

As each thread has a different priority, the question is how these pending thread priorities on a semaphore are handled. To have a more clear view on our test, you can take a look on the expanded diagrams during a small time frame (e.g. one test loop):

- We create 128 threads with different priorities. The creating thread has a lower priority than the threads being created.
- When the thread starts execution, it tries to acquire the semaphore; but as it is taken, the thread stops and the kernel switch back to the creating thread. The time from the acquisition try (which fails) until the creating thread is activated again is called here the "acquisition time". Thus, this time includes the thread switch time. Thread creation takes some time, so the time between each measurement point is large compared with most other tests.
- After the last thread is created and is blocked on the semaphore, the creating thread starts to release the semaphore and this the same number of times as there are blocked threads.
- We start timing at the moment the semaphore is released which in turn will activate the pending thread with the highest priority, which will stop the timing (thus again the thread switch time is included).

Now, the most important part of this test is to see if the number of threads pending on a semaphore has an impact on release times. Clearly, it doesn't, so this is good.

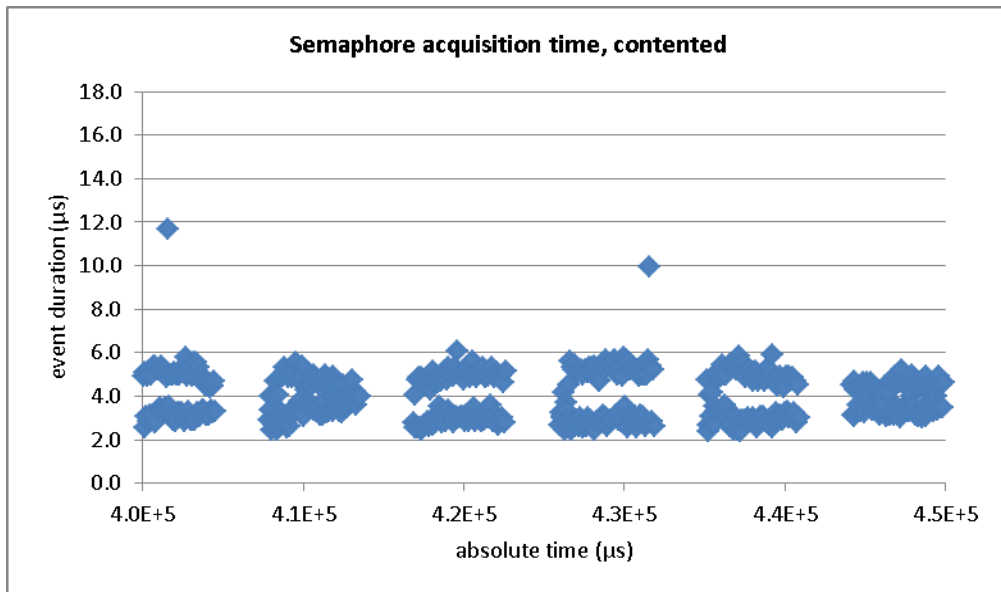
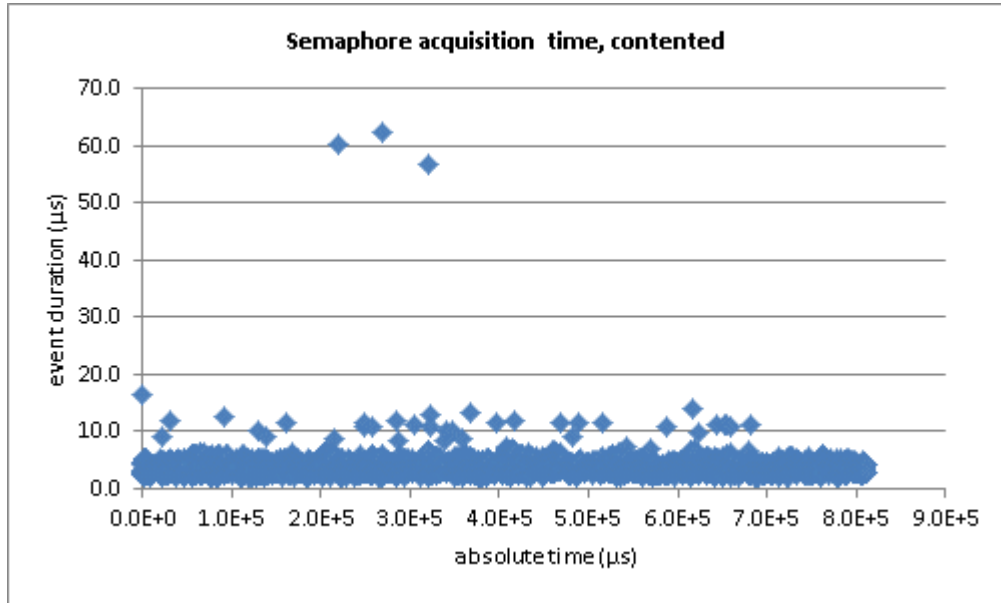
As usual, we find the clock tick back in these diagrams.

4.4.5.1 Test results

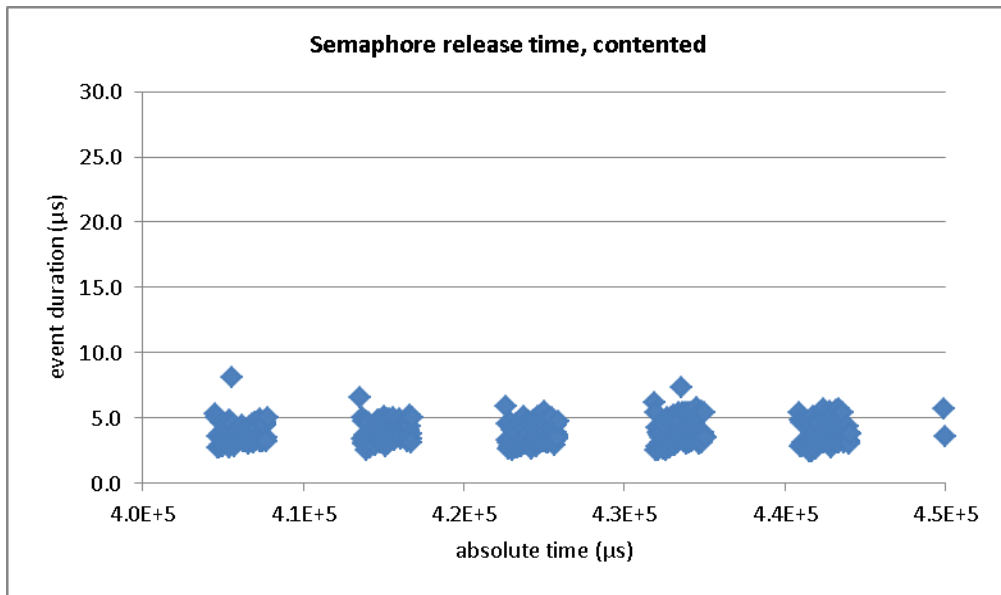
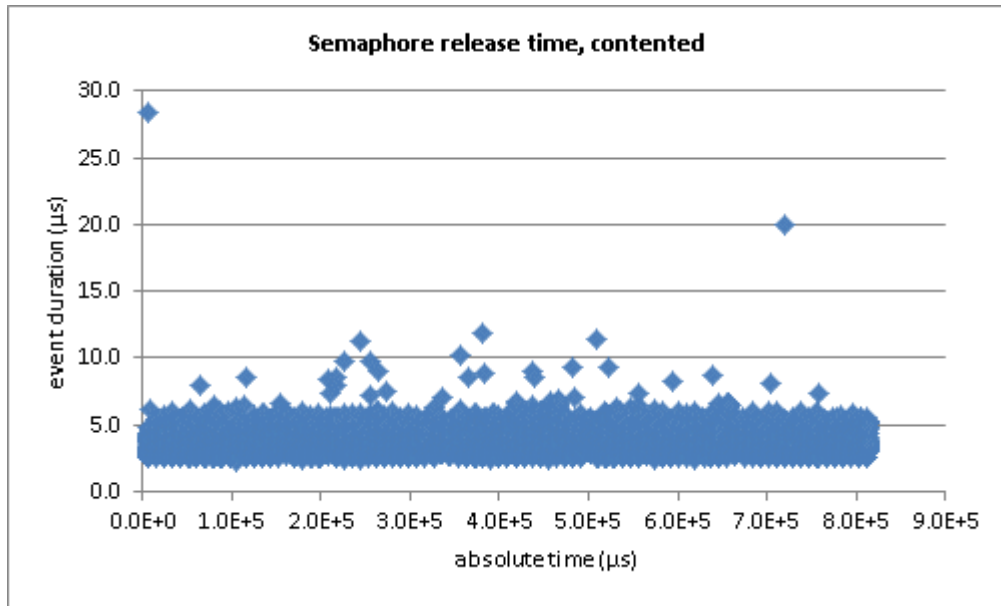
Test	result
Test succeeded	YES
Max number of threads pending	128

Test	Sample qty	Avg	Max	Min
Semaphore acquisition time, contented	7921	3.8 μ s	62.3 μ s	2.2 μ s
Semaphore release time, contented	7921	3.7 μ s	28.5 μ s	2.3 μ s

4.4.5.2 Diagrams



Detailed extract of above



Detailed extract from above

4.5 Mutex tests (MUT)

Here we are going to test the performance and behaviour of the mutual exclusive semaphore.

Although the mutual exclusive semaphore (further called mutex) is mostly explained as being the same as a counting semaphore where the count is one, this is not true. A mutex has a totally different behaviour than semaphores. Mutexes have the concept of "lock owner", and thus they can be used for preventing priority inversions. This is something that cannot be done by semaphores. Therefore it is a bad idea to use semaphores as a critical section protection mechanism.

In scope of the framework, this test will look into detail of a mutex system object that avoids priority inversion.

4.5.1 Priority inversion avoidance mechanism (MUT-B-ARC)

This test will determine if the system call under testing prevents the priority inversion case. Therefore the test will artificially create a priority inversion.

Priority inversion is prevented as expected.

4.5.1.1 Test results

Test	result
Priority inversion avoidance system call present	Yes
System call used	pthread_mutex_lock
Test succeeded	YES
Priority inversion avoided	YES
Mechanism used if any?	pthread_mutexattr_setprotocol: PTHREAD_PRIO_INHERIT

4.5.2 Mutex acquire-release timings: contention case (MUT-P-ARC)

This is the same test as above, but performed in a loop. In this case, the time is measured to acquire and release the mutex in the priority inversion case.

Remark that the acquisition enforces a thread switch. The acquiring thread is blocked and the one having the lock is released. The time is measured from the request for the mutex acquisition to the lower priority thread having the lock being activated.

Before the release, an intermediate priority level thread is activated (between the low priority one having the lock and the high priority one asking the lock). Due to the priority inheritance, this thread does not start to run (the low priority thread having the lock inherited the high priority of the thread asking the lock).

The release time is measured from the release call until the thread requesting the mutex is activated, so it also includes a thread switch.

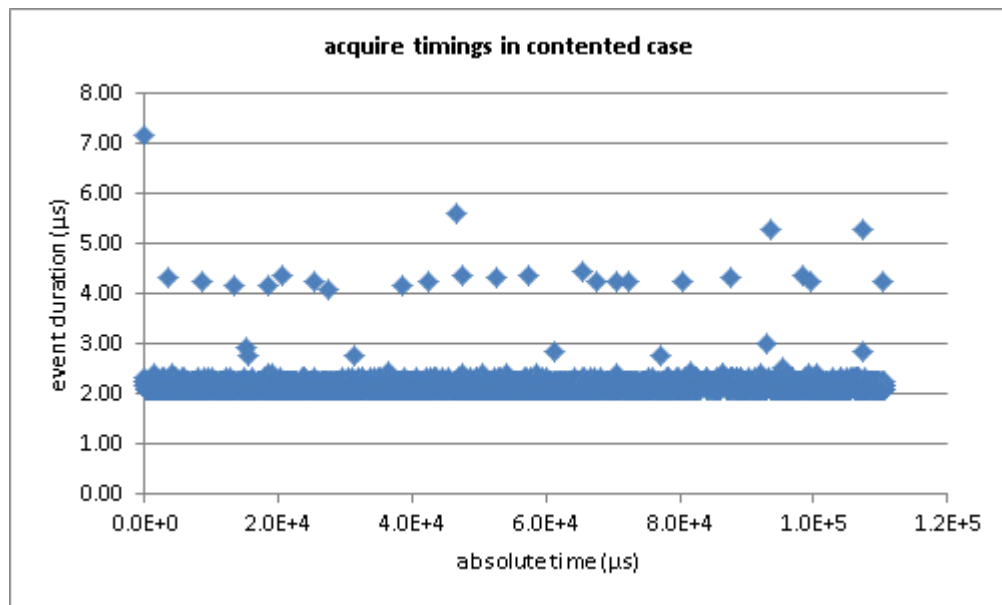
As usual, the clock ticks can be clearly seen.

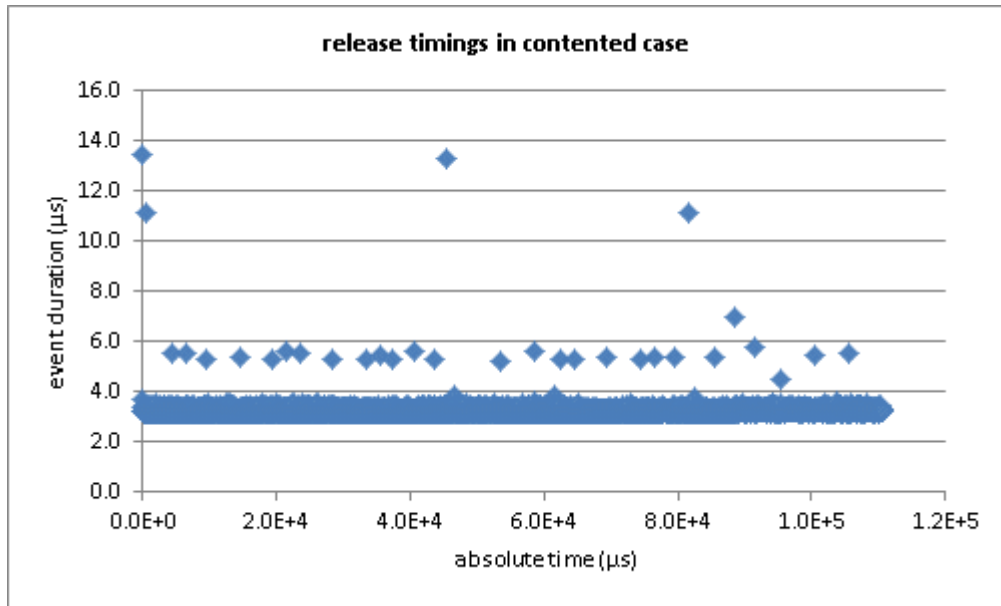
4.5.2.1 Test results

Test	result
Test succeeded	Yes

Test	Sample qty	Avg	Max	Min
Mutex acquisition time, contention	8000	2.2 μ s	7.2 μ s	2.1 μ s
Mutex release time, contention	8000	3.3 μ s	13.5 μ s	3.2 μ s

4.5.2.2 Diagrams:





4.5.3 Mutex acquire-release timings: no-contention case (MUT-P-ARN)

This test measures the overhead of using a lock when it is not locked by another thread. Good designed software will use non-contended locks most of the time and only in some rare cases the lock will be taken by another thread.

Therefore, it is important that the non-contention case should be fast. Remark that this is possible only if the CPU supports some type of atomic instruction, so that no system call is needed when no contention is detected. Clearly, this is the case for QNX; probably no system call is issued. Anyhow, the time it requires starts to be too small to measure.

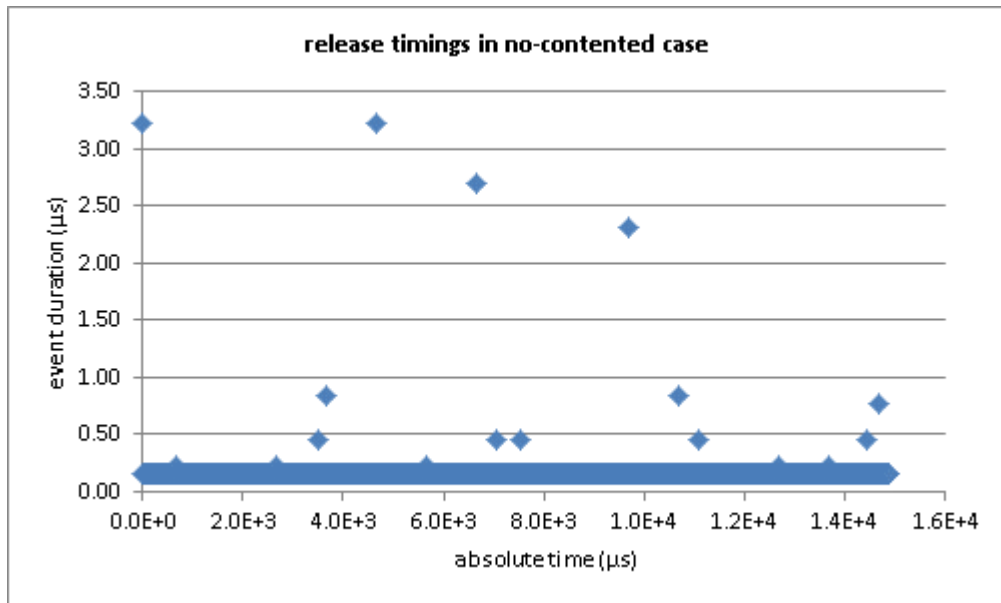
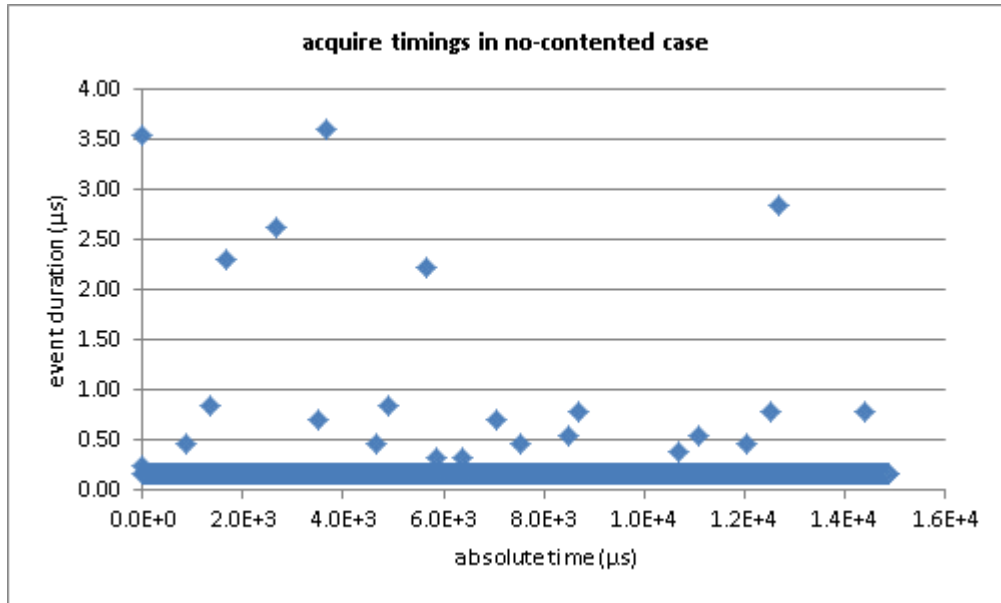
As in all diagrams, the clock tick shows up again.

4.5.3.1 Test results

Test	result
Test succeeded	Yes

Test	Sample qty	Avg	Max	Min
Semaphore acquisition time, no contention	8000	0.2 µs	3.6 µs	0.2 µs
Semaphore release time, no contention	8000	0.2 µs	3.2 µs	0.2 µs

4.5.3.2 Diagrams:



4.6 Interrupt tests (IRQ)

The performance of the interrupt handling in the operating system and hardware is tested here.

In a real-time system, interrupt handling is a major part of the system. Indeed, such systems are typically event driven.

For these tests, interrupts are generated by another General Purpose Timer in the chip (we use already one for tracing purposes). This timer has an independent programmable wrap around timer, so it is not influenced by the RTOS clock. As such, we can guarantee that an independent interrupt source is not synchronised in any way with the platform under test.

4.6.1 Interrupt latency (IRQ_P_LAT)

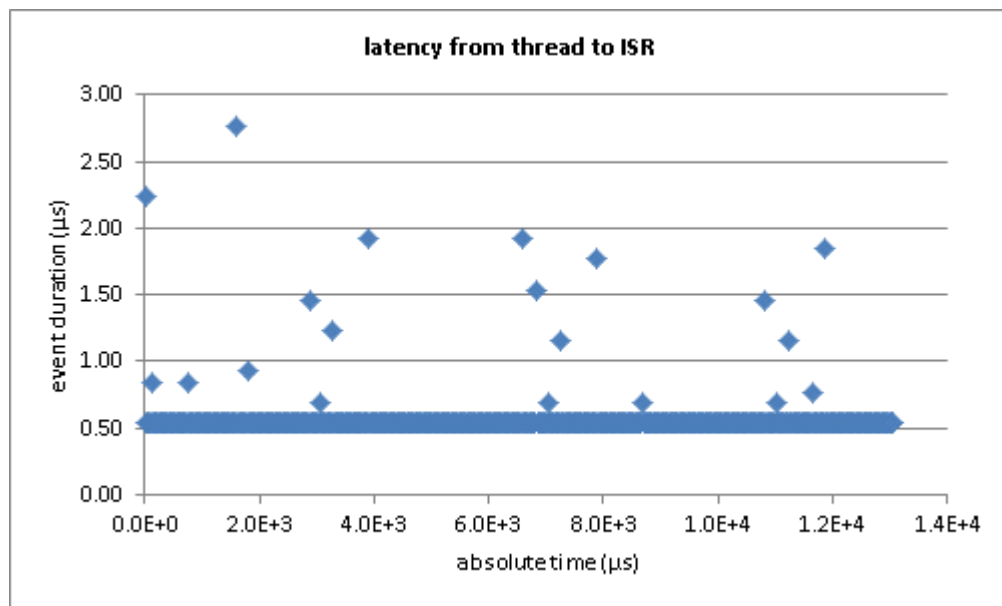
This test measures the time it takes to switch from a running thread to an interrupt handler. The time is measured from the moment the running thread is interrupted which means that it does not measure the hardware interrupt latency.

The results are very good.

4.6.1.1 Test results

Test	Sample qty	Avg	Max	Min
Interrupt dispatch latency	664	0.5 μ s	2.6 μ s	0.5 μ s

4.6.1.2 Diagrams



4.6.2 Interrupt dispatch latency (IRQ_P_DLT)

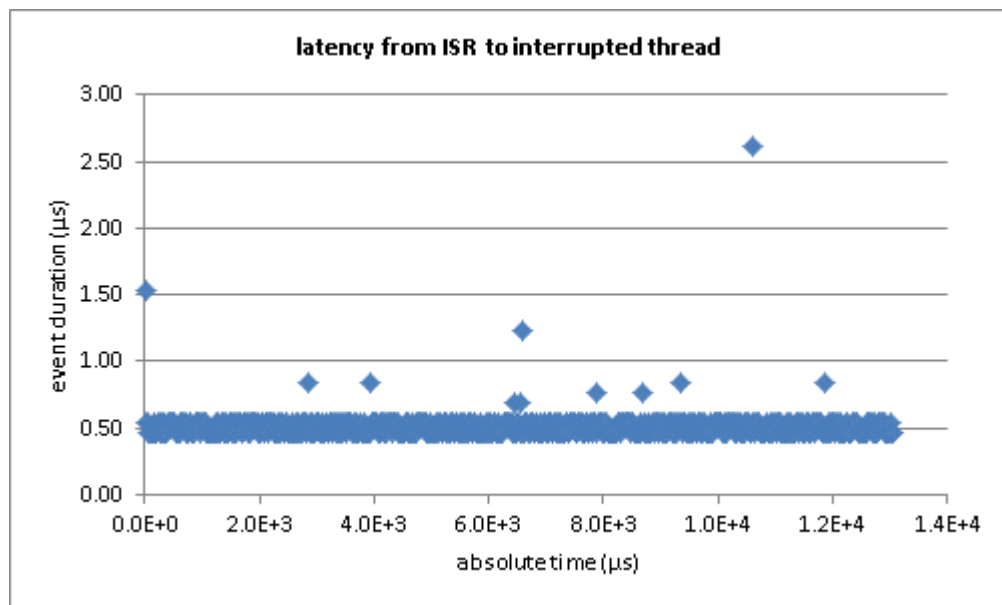
This test measures the time it takes to switch from the interrupt handler back to the interrupted thread.

These results are very good and even better than the (faster) Atom platform; in this case, the measurements can be compared. Thus, although the ARM platform processing power is about half the processing power of the Atom platform, in this scenario it is almost twice faster!

4.6.2.1 Test results

Test	Sample qty	Avg	Max	Min
Dispatch latency from interrupt handler	664	0.6 μ s	2.6 μ s	0.5 μ s

4.6.2.2 Diagrams



4.6.3 Interrupt to thread latency (IRQ_P_TLT)

4.6.3.1 Test results

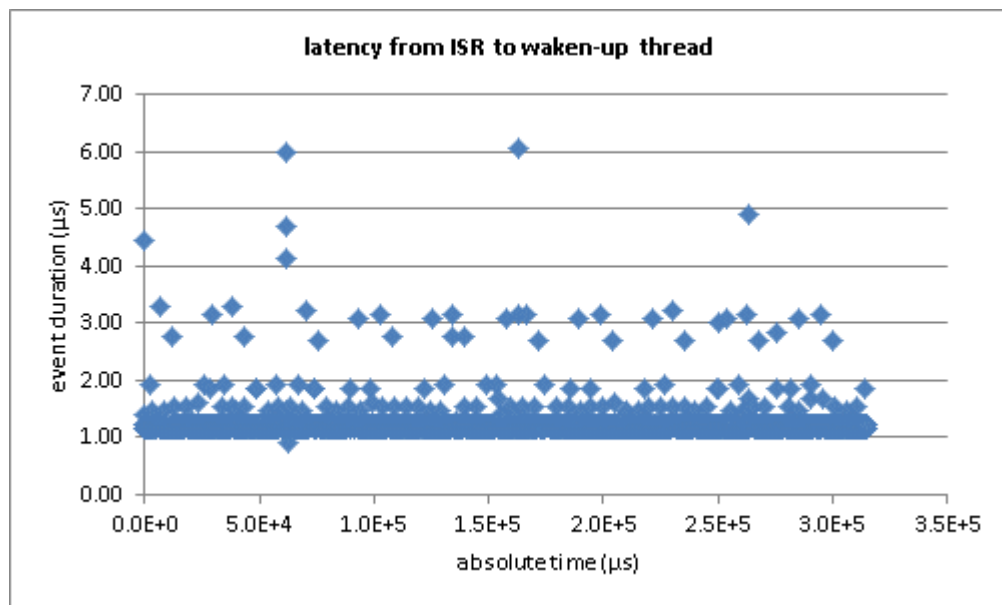
This test measures the time it takes to switch from the interrupt handler to the thread that is activated from the interrupt handler.

This test is done by allowing the interrupt handler to emit an event which releases a blocked thread. This blocking thread has the highest priority in the system. There is also a low priority thread looping. So the measurement takes the time from the interrupt handler to the blocked thread (as a consequence this includes a thread switch).

4.6.3.2 Test results

Test	Sample qty	Avg	Max	Min
Latency from ISR to waken-up thread	16000	1.2 μ s	6.0 μ s	0.9 μ s

4.6.3.3 Diagrams



4.6.4 Maximum sustained interrupt frequency (IRQ_S_SUS)

This test measures the probability that an interrupt is missed. Is the interrupt handling duration stable and predictable?

The test is done on three levels:

- 1000 interrupts, initial phase: a fast test just to see where we have to start searching.
- 1 000 000 interrupts, second phase based on the results from the first phase. This test still takes less than a minute and gives already accurate results.
- 1 000 000 000 interrupts, takes more than 24 hours: to verify stability, therefore we cannot run a lot of tests, especially when it comes to large interrupt latencies.

As is a bit expected knowing the serious difference between cached and un-cached performance, the difference between the best and worst case is significant. At the end, the worst case is about similar like a Pentium MMX 200MHz!

4.6.4.1 Test results

Interrupt period	#interrupts generated	#interrupts serviced	#interrupts lost
9 us	1000	1000	0
6 us	1000	999	1
4 us	1000	996	4
10 us	1000 000	1000 000	0
9 us	1000 000	999 995	5
23 us	1000 000 000	1000 000 000	0
20 us	1000 000 000	999 999 985	54
15 us	1000 000 000	999 999 967	33
10 us	1000 000 000	999 999 991	90

4.7 Memory tests

This test examines the memory leaks of OS.

4.7.1 Memory leak test (MEM_B_LEK)

OS objects are often in a separate pool of memory. Detecting leaks by variously mixing calls to the system APIs will determine if there are some problems that may occur only after long runs.

This test continuously create/remove objects in the operating system (threads, semaphores, mutexes ...).

Test	result
Test succeeded	YES
Test duration (how long we let the endless loop run)	>10h
Number of main test loops done	> 50 000

5 Appendix A: Vendor comments

All vendor comments were integrated within the document as there were no disagreements.

6 Appendix B: Acronyms

Acronym	Explanation
API	Application Programmers Interface: calls used to call code from a library or system.
BSP	Board Support Package: all code and device drivers to get the OS running on a certain board
DSP	Digital Signal Processor
FIFO	First In First Out: a queuing rule
GPOS	General Purpose Operating System
GUI	Graphical User Interface
IDE	Integrated Development Environment (GUI tool used to develop and debug applications)
IRQ	Interrupt Request
ISR	Interrupt Servicing Routine
MMU	Memory Management Unit
OS	Operating System
PCI	Peripheral Component Interconnect: bus to connect devices, used in all PCs!
PIC	Programmable Interrupt Controller
PMC	PCI Mezzanine Card
PrPMC	Processor PMC: a PMC with the processor
RTOS	Real-Time Operating System
SDK	Software Development Kit
SoC	System on a Chip